

Translating Hierarchical Block Diagrams into Composite Predicate Transformers^{*}

Iulia Dragomir¹, Viorel Preoteasa¹, and Stavros Tripakis^{1,2}

¹ Aalto University, Finland

{iulia.dragomir, viorel.preoteasa, stavros.tripakis}@aalto.fi

² University of California, Berkeley, USA

Abstract. Simulink is the de facto industrial standard for designing embedded control systems. When dealing with the formal verification of Simulink models, we face the problem of translating the graphical language of Simulink, namely, hierarchical block diagrams (HBDs), into a formalism suitable for verification. In this paper, we study the translation of HBDs into the compositional refinement calculus framework for reactive systems. Specifically, we consider as target language an algebra of atomic predicate transformers to capture basic Simulink blocks (both stateless and stateful), composed in series, in parallel, and in feedback. For a given HBD, there are many possible ways to translate it into a term in this algebra, with different tradeoffs. We explore these tradeoffs, and present three translation algorithms. We report on a prototype implementation of these algorithms in a tool that translates Simulink models into algebra terms implemented in the Isabelle theorem prover. We test our tool on several case studies including a benchmark Simulink model by Toyota. We compare the three translation algorithms, with respect to size and readability of generated terms, simplifiability of the corresponding formulas, and other metrics.

1 Introduction

Simulink³ is a widely used tool for modeling and simulating embedded control systems. Simulink uses a graphical language based on *hierarchical block diagrams* (HBDs). HBDs are networks of interconnected *blocks*, which can be either *basic* blocks from Simulink’s libraries, or *composite* blocks (*subsystems*), which are themselves HBDs. Hierarchy can be seen as the primary *modularization* mechanism that Simulink offers, which allows to master complexity of large models, improve their readability, and so on.

Our work seeks to develop methods and tools for *compositional* analysis of HBDs, including but not limited to Simulink models. By “compositional” we mean exploiting the hierarchical structure of these diagrams, for instance, reasoning about individual blocks and subsystems independently, and then composing the results to reason about more complex systems. By “analysis”, we mean different types of checks, including exhaustive verification (model-checking), but also more “lightweight” analyses such as *compatibility checking*, which aims to check whether the connections between two or more blocks in the diagram are valid, i.e., are the blocks compatible.

We base our work on the *refinement calculus for reactive systems* [22,19]. In this framework, systems are modeled as *predicate* and *property transformers* [2,19]. Open, non-deterministic, and non-input-receptive systems can be modeled in the framework. *Serial*, *parallel*, and *feedback* composition operators can be used to form more complex systems from simpler ones. Compatibility can be checked during such compositions. Both safety and liveness properties can be expressed in the framework (e.g., using LTL) and used for verification. In addition, the framework includes the notion of *refinement*, a binary relation between components, which characterizes *substitutability* (when can a component replace another one while preserving system properties). Refinement has multiple usages, including compositional and incremental design, and reusability. This makes the framework compelling for application on tools like Simulink, which have a naturally compositional hierarchical language.

^{*} This work has been partially supported by the Academy of Finland, the U.S. National Science Foundation (awards #1329759 and #1139138), and by UC Berkeley’s iCyPhy Research Center (supported by IBM and United Technologies).

³ <http://www.mathworks.com/products/simulink/>




Fig. 1: Three ways to view and translate the same block diagram.

In order to use refinement calculus⁴ with Simulink, we need to translate the graphical language of Simulink (HBDs) into the algebra of composite transformers. This translation raises interesting problems, and these are the topic of this paper.

To illustrate some of the questions that arise, consider the block diagram in Fig. 1a. Let P_A and P_B be transformers modeling the blocks A and B in the diagram. How should we compose P_A and P_B in order to get a transformer that represents the entire diagram? As it turns out, there are many such compositions possible. One option is to compose first P_A and P_B in series, and then compose the result in feedback, following Fig. 1a. This results in the composite transformer $\text{feedback}_a(P_A \circ (P_B \parallel \text{Id}))$, where \circ is composition in series, \parallel in parallel, and feedback_x is feedback applied on port x . Id is the transformer representing the identity function. A has two outputs and B only one input, therefore to connect them in series we first form the parallel composition $P_B \parallel \text{Id}$, which represents a system with two inputs.

Another option is to compose the blocks in series in the opposite order, P_B followed by P_A , and then apply feedback. This results in the transformer $\text{feedback}_c((P_B \parallel \text{Id}) \circ P_A)$. A third option is to compose the two blocks first in parallel, and then apply feedback on the two ports a, c . This results in the transformer $\text{feedback}_{a,c}(P_A \parallel P_B)$. Although semantically equivalent, these three transformers have different computational properties.

Clearly, for complex diagrams, there are many possible translation options. In this paper we study these options in depth. Our main contributions are the following. First, we present three different translation strategies: *feedback-parallel* translation which forms the parallel composition of all blocks, and then applies feedback; *incremental* translation which orders blocks topologically and composes them one by one; and *feedbackless* translation, which avoids feedback composition altogether, provided the original block diagram has no algebraic loops. Second, we discuss the tradeoffs of these strategies, in terms of several metrics, including the size of the resulting expressions, readability, and also *simplifiability*, which can generally be defined as how easy/automatic it is to simplify the expressions into simple formulas characterizing the overall system. Third, we report on a prototype tool which implements the three translation strategies. The tool takes as input hierarchical Simulink models and generates composite predicate transformers implemented in the Isabelle proof assistant⁵. Isabelle or other tools are then used to simplify the corresponding formulas, to perform verification, etc. We evaluate and compare the three translation strategies on several case studies, including a Fuel Control System benchmark by Toyota [11,12].

Paper structure. In §2 we present the basic modeling notions in Simulink for designing hierarchical systems and their semantics. We illustrate these notions on a toy running example, a 1-step counter. §3 depicts the predicate transformer algebra containing atomic and composite terms and three composition operators. Some examples on how to formalize atomic Simulink blocks with predicate transformers are presented. In §4, the three translation strategies of Simulink hierarchical block diagrams to composite predicate transformers are described and illustrated on examples. We report on the tool support in §5 and we validate the approach on two case studies. We discuss the benefits and drawbacks of each translation strategy with respect to different metrics. Finally, we discuss the related work before concluding.

⁴ but also any other framework that relies on an algebra of atomic components composed with operators such as serial, parallel, and feedback

⁵ <https://isabelle.in.tum.de/>




Fig. 2: Simulink model of a counter with step 1.

2 Hierarchical Block Diagrams in Simulink

A Simulink HBD is a network of blocks interconnected by wires. Blocks can be either *basic* blocks from the Simulink libraries, or *composite* blocks (*subsystems*). A basic block is described by: (1) a label, (2) a list of parameters, (3) a list of in- and out-ports, (4) a vector of state variables with predefined initial values (i.e., the local memory of a block) and (5) functions to compute the outputs and next state variables. The outputs are computed from the inputs, current state and parameters. State variables are updated by a function with the same arguments. Subsystems are defined by their label, list of in- and out-ports, and the list of block instances that they contain – both atomic and composite.

Simulink allows to model both discrete and continuous-time blocks. For example, `UnitDelay` is a discrete-time block which outputs at step n the input at step $n - 1$. An `Integrator` is a continuous-time block whose output is described by a differential equation solved with numerical methods. We interpret a Simulink model as a discrete-time model (essentially an input-output state machine, possibly infinite-state) which evolves in a sequence of discrete steps. Each step has duration Δt , which is a parameter (user-defined or automatically computed by Simulink based on the blocks' time rates). In this paper we consider single-rate models. We also assume that Simulink models are free from *algebraic loops*, that is, feedback loops resulting in instantaneous cyclic dependencies. We can have feedback loops, but they must be “broken” by blocks such as `UnitDelay` or `Integrator`, as in the example that follows.

Running example. Throughout the paper we illustrate our methods using a simple example of a counter, shown in Fig. 2. This is a hierarchical (two-level) Simulink model. The top-level diagram (Fig. 2a) contains three block instances: the step of the counter as a `Constant` basic block, the subsystem `DelaySum`, and the `Scope` basic block which allows to view simulation results. The subsystem `DelaySum` (Fig. 2b) contains a `UnitDelay` block instance which represents the state of the counter. `UnitDelay` can be specified by the formula $a = s \wedge s' = c$, where c is the input, a the output, s the current state and s' the next state variable. We assume that s is initially 0. The `Add` block instance adds the two input values and outputs the result in the same time step: $c = f + e$. The *junction* after wire a (black dot in the figure) can be seen as a basic block duplicating (or *splitting*) its input to its two outputs: $f = a \wedge g = a$.

3 Atomic and Composite Predicate Transformers

Monotonic predicate transformers [7] (MPTs) are an expressive formalism, used within the context of programming languages to model non-determinism, correctness (both functional correctness and termination), and refinement [2]. In this section we show how MPTs can also model input-output systems such as state machines, and by extension also capture the semantics of block diagrams *à la* Simulink. The idea is to represent the computation step of a state machine using an MPT which has the inputs of the state machine and the *current state* variable as inputs, and the outputs of the state machine and the *next state* variable as outputs. We also discuss MPT composition operators.

3.1 Monotonic Predicate Transformers for Modeling Systems

A *predicate* on Σ is a function $q : \Sigma \rightarrow \text{Bool}$. Predicate q can also be seen as a subset of Σ : for $\sigma \in \Sigma$, σ belongs to the subset iff $q(\sigma)$ is true. Predicates can be ordered by the subset relation: we write $q \leq q'$ if predicate q , viewed as a set, is a subset of q' . $\text{Pred}(\Sigma)$ denotes the set of predicates $\Sigma \rightarrow \text{Bool}$.

A *predicate transformer* is a function $S : (\Sigma' \rightarrow \text{Bool}) \rightarrow (\Sigma \rightarrow \text{Bool})$, or equivalently, $S : \text{Pred}(\Sigma') \rightarrow \text{Pred}(\Sigma)$. S takes a predicate on Σ' and returns a predicate on Σ . S is *monotonic* if $\forall q, q' : q \leq q' \Rightarrow S(q) \leq S(q')$.

Traditionally, MPTs have been used to model sequential programs using weakest precondition semantics. Given a MPT $S : (\Sigma' \rightarrow \text{Bool}) \rightarrow (\Sigma \rightarrow \text{Bool})$, and a predicate $q' : \Sigma' \rightarrow \text{Bool}$ capturing a set of *final* states, $S(q')$ captures the set of all *initial* states, such that if the program is started in any state in $S(q')$, it is guaranteed to finish in some state in q' . But this is not the only possible interpretation of S . S can also model input-output systems. For instance, S can model a *stateless* system with a single inport ranging over Σ , and a single outport ranging over Σ' . Given a predicate q' characterizing a set of possible *output values*, $S(q')$ characterizes the set of all *input values* which, when fed into the system, result in the system outputting a value in q' . As an example, the identity function can be modeled by the MPT $\text{Id} : \text{Pred}(\Sigma) \rightarrow \text{Pred}(\Sigma)$, defined by $\text{Id}(q) = q$, for any q .

MPTs can also model *stateful* systems. For instance, consider the **UnitDelay** described in §2. Let the input, output, and state variable of this system range over some domain Σ . Then, this system can be modeled as a MPT $S : \text{Pred}(\Sigma \times \Sigma) \rightarrow \text{Pred}(\Sigma \times \Sigma)$. The Cartesian product $\Sigma \times \Sigma$ captures pairs of (*input, current state*) or (*output, next state*) values. Intuitively, we can think of this system as a function which takes as input (x, s) , the input x and the current state s , and returns (y, s') , the output and the next state s' , such that $y = s$ and $s' = x$. The MPT S can then be defined as follows:

$$S(q) = \{(x, s) \mid (s, x) \in q\}.$$

In the definition above we view predicates q and $S(q)$ as sets.

Syntactically, a convenient way to specify systems is using formulas on input, output, and state variables. For example, the identity system can be specified by the formula $y = x$, where y is the output variable and x is the input. The **UnitDelay** system can be specified by the formula $y = s \wedge s' = x$. We next introduce operators which define MPTs from predicates and relations.

For a predicate $p : \Sigma \rightarrow \text{Bool}$ and a relation $r : \Sigma \rightarrow \Sigma' \rightarrow \text{Bool}$, we define the *assert* MPT, $\{p\} : \text{Pred}(\Sigma) \rightarrow \text{Pred}(\Sigma)$, and the *non-deterministic update* MPT, $[r] : \text{Pred}(\Sigma') \rightarrow \text{Pred}(\Sigma)$, where:

$$\{p\}(q) = (p \wedge q) \quad \text{and} \quad [r](q) = \{\sigma \mid \forall \sigma' : \sigma' \in r(\sigma) \Rightarrow \sigma' \in q\}$$

Transformer $\{p\}$ is used to model non-input-receptive systems, that is, systems where some inputs are illegal [22]. $\{p\}$ constrains the inputs so that they must satisfy predicate p . It accepts only those inputs and behaves like the identity function. That is, $\{p\}$ models a partial identity function, restricted to the domain p . Transformer $[r]$ models an input-receptive but possibly non-deterministic system. Given input σ , the system chooses non-deterministically some output σ' such that $\sigma' \in r(\sigma)$ is true. If no such σ' exists, then the system behaves *miraculously* [2]. In our framework we ensure non-miraculous behavior, as explained below.

To model basic Simulink blocks, we often combine $\{p\}$ and $[r]$ using the *serial composition* operator \circ , which for predicate transformers is simply function composition. Given two MPTs $S : \text{Pred}(\Sigma_2) \rightarrow \text{Pred}(\Sigma_1)$ and $T : \text{Pred}(\Sigma_3) \rightarrow \text{Pred}(\Sigma_2)$, their serial composition $(S \circ T) : \text{Pred}(\Sigma_3) \rightarrow \text{Pred}(\Sigma_1)$ is defined as $(S \circ T)(q) = S(T(q))$.

For example, consider a system with two inputs x, y and one output z , performing the division $z = \frac{x}{y}$. We want to state that division by zero is illegal, and therefore, the system should reject any input where $y = 0$. This system can be specified as the MPT

$$\text{Div} = \{\lambda x, y : y \neq 0\} \circ [\lambda(x, y), z : z = \frac{x}{y}]$$

where we employ lambda-notation for functions. Note that $[\lambda(x, y), z : z = \frac{x}{y}]$ alone is not enough to capture **Div**, because it allows illegal inputs where $y = 0$ (and then behaves miraculously). The same is true for $[\lambda(x, y), z : y \neq 0 \wedge z = \frac{x}{y}]$. To ensure non-miraculous behavior, we always model non-input-receptive systems using a suitable assert transformer $\{p\}$.

For a function $f : \Sigma \rightarrow \Sigma'$ the *functional update* $[f] : \text{Pred}(\Sigma') \rightarrow \text{Pred}(\Sigma)$ is defined as $[\lambda\sigma, \sigma' : \sigma' = f(\sigma)]$ and we have

$$[f](q) = \{\sigma \mid f(\sigma) \in q\} = f^{-1}(q)$$

Functional predicate transformers are of the form $\{p\} \circ [f]$, and *relational predicate transformers* are of the form $\{p\} \circ [r]$, where p is a predicate, f is a function, and r is a relation. *Atomic predicate transformers*

are either functional or relational transformers. Div is a functional predicate transformer which can also be written as $\text{Div} = \{\lambda x, y : y \neq 0\} \circ [\lambda x, y : \frac{x}{y}]$.

For assert and update transformers based on Boolean expressions we introduce a simplified notation that avoids lambda abstractions. If P is Boolean expression on some variables x_1, \dots, x_n , then $\{x_1, \dots, x_n : P\}$ denotes the assert transformer $\{\lambda x_1, \dots, x_n : P\}$. Similarly if R is a Boolean expression on variables $x_1, \dots, x_n, y_1, \dots, y_k$ and F is a tuple of numerical expressions on variables x_1, \dots, x_n , then $[x_1, \dots, x_n \rightsquigarrow y_1, \dots, y_k : R]$ and $[x_1, \dots, x_n \rightsquigarrow F]$ are notations for $[\lambda(x_1, \dots, x_n), (y_1, \dots, y_k) : R]$ and $[\lambda x_1, \dots, x_n : F]$, respectively. With these notations the Div transformer becomes:

$$\text{Div} = \{x, y : y \neq 0\} \circ [x, y \rightsquigarrow \frac{x}{y}]$$

3.2 More Examples of Atomic Predicate Transformers

We already saw how to model basic Simulink blocks like Id and Div. Next we provide more examples, such as constants, delays, and integrators. A constant block parameterized by constant c has no input, and a single output equal to c . As a predicate transformer the constant block has as input the empty tuple $()$, and outputs the constant c :

$$\text{Const}(c) = [() \rightsquigarrow c]$$

The unit delay block is modeled as the atomic predicate transformer

$$\text{UnitDelay} = [x, s \rightsquigarrow s, x]$$

Simulink includes continuous-time blocks such as the *integrator*, which computes the integral $\int_0^x f$ of a function f . Simulink uses different integration methods to simulate this block. We use the Euler method with fixed time interval dt , a parameter. If x is the input, y the output, and s the state variable of the integrator, then $y = s$ and $s' = s + x \cdot dt$. Therefore, the integrator can be modeled as the MPT

$$\text{Integrator}(dt) = [x, s \rightsquigarrow s, s + x \cdot dt]$$

All other Simulink atomic blocks fall within these cases discussed above. Relation (1) introduces the definitions of some blocks that we use in our examples.

$$\text{Add} = [x, y \rightsquigarrow x + y] \quad \text{Split} = [x \rightsquigarrow x, x] \quad \text{Scope} = \text{Id} \quad (1)$$

3.3 Composite Predicate Transformers

Basic Simulink blocks are modeled using atomic predicate transformers. To model arbitrary block diagrams, we use *composite predicate transformers* (CPTs). These are expressions over the atomic predicate transformers using serial, *parallel*, and *feedback* composition operators. Serial composition \circ was already introduced in §3.1. Next we define the parallel and the feedback operators.

For $S : \text{Pred}(Y) \rightarrow \text{Pred}(X)$, and $T : \text{Pred}(Y') \rightarrow \text{Pred}(X')$, the parallel composition of S and T , denoted $S \parallel T : \text{Pred}(Y \times Y') \rightarrow \text{Pred}(X \times X')$ is defined as

$$(S \parallel T)(q) = \{(x, x') \mid \exists p, p' : p \times p' \leq q \wedge x \in S(p) \wedge x' \in T(p')\}$$

For $S : \text{Pred}(U \times Y) \rightarrow \text{Pred}(U \times X)$, the feedback of S , denoted $\text{feedback}(S) : \text{Pred}(Y) \rightarrow \text{Pred}(X)$ is defined as

$$\text{feedback}(S) = [x \rightsquigarrow x, x] \circ (\text{sel}(S) \parallel \text{Id}) \circ S \circ [v, y \rightsquigarrow y]$$

where

$$\text{sel}(S) = \{x : \exists u : (u, x) \in S(\top)\} \circ [x \rightsquigarrow u, x : (u, x) \in S(\top)] \circ S \circ [v, y \rightsquigarrow v]$$

In this definition \top is the total predicate ($\top.x$ is true for all x), and $S(\top)$ is the set of all legal inputs of S . The feedback operation is designed such that it provides the expected result for a PT S when its first output v depends only on the second input x . Intuitively this case is represented in Fig. 3c. If S has this special form, then $\text{sel}(S)$ selects the component T' ($\text{sel}(S) = T'$). The component T is $S \circ [v, y \rightsquigarrow y]$, and the feedback of S is represented in Figure 3d: $[x \rightsquigarrow x, x] \circ (T' \parallel \text{Id}) \circ T$. Proper Simulink diagrams (those without algebraic loops) satisfy this property.

Next we introduce two lemmas for unfolding the composition operators when they are applied to relational and functional predicate transformers. For serial and parallel compositions we have:




Fig. 3: (a) MPT S , (b) $\text{feedback}(S)$, (c) special case of S , (d) feedback of (c).

Lemma 1. For p, p', r, r', f , and f' of appropriate types we have:

- (i) $\{p\} \circ [r] \circ \{p'\} \circ [r'] = \{x : x \in p \wedge (\forall y : y \in r(x) \Rightarrow y \in p')\} \circ [r \circ r']$
- (ii) $\{p\} \circ [f] \circ \{p'\} \circ [f'] = \{p \wedge (p' \circ f)\} \circ [f' \circ f]$
- (iii) $(\{p\} \circ [r]) \parallel (\{p'\} \circ [r']) = \{x, y : x \in p \wedge y \in p'\} \circ [x, y \rightsquigarrow u, v : u \in r(x) \wedge v \in r'(y)]$
- (iv) $(\{p\} \circ [f]) \parallel (\{p'\} \circ [f']) = \{x, y : x \in p \wedge y \in p'\} \circ [x, y \rightsquigarrow f(x), f'(y)]$

For feedback we have:

Lemma 2. For p, r, f and f' of appropriate types we have:

- (i) $\text{feedback}(\{p\} \circ [r]) =$
 $\{x : (\exists u : (u, x) \in p) \wedge (\forall a : (\exists u : (u, x) \in p \wedge (\exists y : (a, y) \in r(u, x))) \Rightarrow (a, x) \in p))\}$
 $\circ [x \rightsquigarrow y : (\exists v : (\exists u : (u, x) \in p \wedge (\exists y : (v, y) \in r(u, x))) \wedge (v, y) \in r(v, x)))]$
- (ii) $\text{feedback}(\{p\} \circ [u, x \rightsquigarrow f'(x), f(u, x)]) = \{x : (f'(x), x) \in p\} \circ [x \rightsquigarrow f(f'(x), x)]$

Lemma 2.(i) gives a general formula for unfolding the feedback of a relational predicate transformer, and Lemma 2.(ii) shows that the feedback works as expected when applied to a functional predicate transformer as in Fig.3c. Although feedback operations applied to a proper Simulink diagrams are always of the form from Fig.3c, when translating arbitrary diagrams *compositionally* to CPTs, we may not be able to identify the components f' and f . Because of this, to simplify the CPTs, we need to use Lemma 2.(i), and rely on a powerful simplification mechanism that will eliminate the quantifiers. Another problem with Lemma 2.(i) is that, although atomic Simulink blocks are functional, after applying the feedback we obtain a relational PT.

As an illustration of how CPTs are used in this work, consider our running example (Fig. 2). An example translation of the DelaySum subsystem and of the top-level Simulink model yields the following two CPTs:

$$\begin{aligned} \text{DelaySum} &= \text{feedback}((\text{Add} \parallel \text{Id}) \circ \text{UnitDelay} \circ (\text{Split} \parallel \text{Id})) \\ \text{Counter} &= (\text{Const}(1) \parallel \text{Id}) \circ \text{DelaySum} \circ (\text{Scope} \parallel \text{Id}) \end{aligned} \quad (2)$$

The Id transformers in these definitions are for propagating the state introduced by the unit delay. Expanding the definitions of the atomic blocks, and applying Lemmas 1.(iv), 2.(ii), and 1.(ii) we obtain the *simplified* definitions:

$$\text{DelaySum} = [x, s \rightsquigarrow s, s + x] \quad \text{and} \quad \text{Counter} = [s \rightsquigarrow s, s + 1] \quad (3)$$

Our goal is to perform automatically, (a) the translation from Simulink diagrams (Fig. 2) to CPTs (Defs.(2)); and (b) the expansion and simplification from (2) to (3). In the next section we discuss problem (a). For (b), we use Isabelle and related tools.

4 From Hierarchical Block Diagrams to Composite Predicate Transformers

We want to translate HBDs to CPTs.⁶ As illustrated in the introduction, this mapping is not unique: for a given HBD, there are many possible CPTs that we could generate. As we shall see in §5, these CPTs, although semantically equivalent, have different simplifiability properties. In this section, we describe three different translation strategies.

The overall translation flow consists in the following steps:

⁶ We mention that this model-to-text translation is generic enough to accept as input any graphical hierarchical block diagrams, and it is not restricted to Simulink.

1. Parsing of the Simulink model diagram(s) in order to extract the relevant information about the system structure (i.e. blocks, connectors) and its processing (e.g., wire instantiation).
2. Mapping of each elementary atomic block to an atomic property transformer. These PTs are generated at runtime based on the customizable number of inputs and outputs of atomic blocks. Examples of PTs for atomic blocks have been provided in §3.1.
3. Building CPTs from hierarchical structures by applying different translation strategies on components' compositions.
4. Producing a complete Isabelle theory subjected to analysis and verification.

This process is automatized by a compiler which is described in §5.1. In the following we focus on the translation strategies for HBDs.

4.1 Internal Representation of HBDs

In order to manipulate Simulink HBDs, we define an internal representation of a Simulink model as a data structure in our tool. This internal representation is faithful to the graphical one. Each block instance A consists of a list of inputs $A.in$, a list of outputs $A.out$ and a predicate transformer $A.cpt$. Moreover, if the block instance corresponds to a subsystem it will also enclose the list of contained components. By *component* we refer in the following to block instances as described by the internal representation.

The lists of inputs/outputs are made of variables (i.e. names), which in case of vector data are indexed over the number of elements in the vector. Wires between block instances are modeled by variables name matching, i.e., the output of a component has the same name as the input of its target. We use the following operations on input/output lists: (1) $+$ is list concatenation, (2) \cap is list intersection which preserves the order of elements from the first operand, and (3) \setminus is list difference which again preserves the order of elements of the first operand. The CPT for a component either is of a predefined type in case of an atomic block or is obtained by applying composition operators on the contained components' CPTs.

In what follows, we describe how a *flat* (non-hierarchical), *connected* diagram is translated. If the diagram consists of many disconnected “islands”, we can simply translate each island separately. Hierarchical diagrams are translated bottom-up: we first translate the subsystems, then their parent, and so on.

4.2 Composition Algorithms for Simulink Components

In the following we design and implement algorithms that apply the composition operators – parallel, serial, and feedback – on any two components A and B for producing the CPT. These algorithms are then used within the translation strategies to obtain the diagram's CPT.

Firstly, we present a decision procedure (Algorithm 1) that determines which composition operators(s) should be applied based on the dependencies between A and B . Also, components are ordered such that the number of “eliminated” wires is maximized. If A and B are not connected, parallel composition is applied. Otherwise, serial composition is used, possibly together with feedback if necessary.

On the example from Fig.1a, $\text{COMPOSE}(A, B)$ will yield the composition of A and B in series encased in a feedback. $\text{COMPOSE}(B, A)$ gives the feedback of the serial composition of B and A , represented in Fig.1b. Fig.1c will not be produced by the decision procedure.

Each of the algorithms **PARALLEL**, **FEEDBACK**, and **SERIAL** creates a new component that replaces within the diagram the components A and B . Therefore, we need to compute the list of inputs, outputs and the CPT.

Algorithm 2 details the component obtained by applying the parallel composition operator between A and B . The lists of inputs/outputs of the new component res are given by the concatenation of the individual lists of inputs and outputs respectively. The CPT is written as $P_A \parallel P_B$. On Fig.1c, **PARALLEL**(A, B) gives a component with inputs $\{a, b, c\}$, outputs $\{c, d, a\}$, and the previous CPT.

The **FEEDBACK** algorithm represented in Algorithm 3 applies on a component A with the feedback variables computed as the intersection between its inputs and outputs. The new component has those inputs and outputs which are not feedback variables. The CPT is written as the successive application of the **feedback** operator on all feedback variables. Recall that the **feedback** operator applies on one variable at a time. The feedback variables must come first in matching order in the input and output

Algorithm 1 Decision procedure for the composition of two components A and B .

```
1: procedure COMPOSE( $A, B$ )
2:    $A2B\_wires \leftarrow A.out \cap B.in$ 
3:    $B2A\_wires \leftarrow B.out \cap A.in$ 
4:   if  $A2B\_wires = \emptyset$  and  $B2A\_wires = \emptyset$  then
5:     return PRODUCT( $A, B$ )
6:   end if
7:   if  $|B2A\_wires| \leq |A2B\_wires|$  then
8:     if  $|B2A\_wires| = 0$  then
9:       return SERIAL( $A, B$ )
10:    else
11:      return FEEDBACK(SERIAL( $A, B$ ))
12:    end if
13:  else
14:    if  $|A2B\_wires| = 0$  then
15:      return SERIAL( $B, A$ )
16:    else
17:      return FEEDBACK(SERIAL( $B, A$ ))
18:    end if
19:  end if
20: end procedure
```

Algorithm 2 Parallel operator-based composition for two blocks A and B .

```
1: PRODUCT( $A, B$ )
2:    $res.in \leftarrow A.in + B.in$ 
3:    $res.out \leftarrow A.out + B.out$ 
4:    $res.cpt \leftarrow A.cpt \parallel B.cpt$ 
5:   return  $res$ 
```

lists. Therefore we need to reorder variables, which is achieved by the first and last expressions of the feedback operator on line 5.

Algorithm 3 Feedback operator application on a block A .

```
1: FEEDBACK( $A$ )
2:    $fdbv \leftarrow A.in \cap A.out$ 
3:    $res.in \leftarrow A.in \setminus fdbv$ 
4:    $res.out \leftarrow A.out \setminus fdbv$ 
5:    $res.cpt \leftarrow [fdbv + res.in \rightsquigarrow A.in] \circ A.cpt \circ [A.out \rightsquigarrow fdbv + res.out]$ 
6:   for  $i=1$  to  $|fdbv|$  do
7:      $res.cpt \leftarrow \text{feedback}(res.cpt)$ 
8:   end for
9:   return  $res$ 
```

On Fig.1c, FEEDBACK(PARALLEL(A, B)) gives the component with input b and output d . The explicit CPT is written: $\text{feedback}(\text{feedback}([a, c, b \rightsquigarrow a, b, c] \circ (P_A \parallel P_B) \circ [c, d, a \rightsquigarrow a, c, d]))$.

The SERIAL function, described by Algorithm 4, proceeds as follows. It computes the unmatched inputs of B and the unmatched outputs of A . The inputs of the new component are those of A and the unmatched inputs of B . The outputs are those of B and the unmatched outputs of A . We transfer the unmatched inputs/outputs with the ld put in parallel with the according component. The the CPT is given as the serial composition of these expressions, interlaced with a variable reordering expression (line 16) to match the components.

On Fig.1a, the algorithm creates a component with inputs $\{a, b\}$ and outputs $\{a, d\}$. Since d is an unmatched output of A , the CPT of B is composed in parallel with ld . Therefore, we obtain the CPT: $P_A \circ (P_B \parallel \text{ld})$. On Fig.1b, SERIAL(B, A) yields a component with inputs $\{c, b\}$ and outputs $\{c, d\}$. Now,

Algorithm 4 Serial operator-based composition of two blocks A and B .

```

1: SERIAL( $A, B$ )
2:    $B\_unc\_in \leftarrow B.in \setminus A.out$ 
3:    $res.in \leftarrow A.in + B.unc\_in$ 
4:    $A\_unc\_out \leftarrow A.out \setminus B.in$ 
5:    $res.out \leftarrow B.out + A\_unc\_out$ 
6:   if  $|B\_unc\_in| = 0$  then
7:      $cptA \leftarrow A.cpt$ 
8:   else
9:      $cptA \leftarrow A.cpt \parallel Id$ 
10:  end if
11:  if  $|A\_unc\_out| = 0$  then
12:     $cptB \leftarrow B.cpt$ 
13:  else
14:     $cptB \leftarrow B.cpt \parallel Id$ 
15:  end if
16:   $res.cpt \leftarrow cptA \circ [A.out + B.unc\_in \rightsquigarrow B.in + A.unc\_out] \circ cptB$ 
17:  return  $res$ 

```

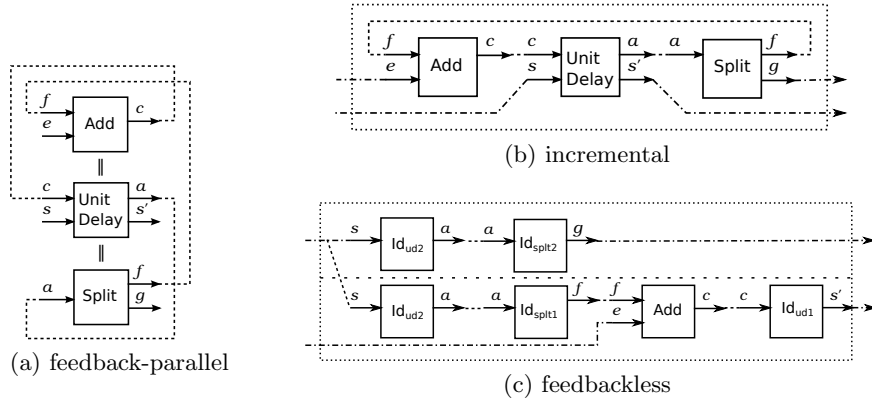


Fig. 4: Translation strategies for the DelaySum subsystem of Fig.2b.

b is an unmatched input for A and again we use Id to transfer it. Then, the CPT is: $(P_B \parallel Id) \circ P_A$. To obtain the diagrams CPTs, we apply Algorithm 3.

4.3 Translation Algorithms for HBDs

Hierarchical block diagrams are translated into CPTs by applying a mixture of the three composition operators (and in consequence algorithms) on the predicate transformers corresponding to the contained components. There are several strategies in which (C)PTs can be assembled together accordingly to the block diagram: (1) by considering all components “running” in parallel and reconnecting them through feedback, (2) exploiting the connections between components and maximizing the usage of serial compositions and (3) reorganizing components (after some syntactic sugar constructions) in order to eliminate the feedback operator and use only the parallel and serial ones. As we will show in §5, a minimal number of applied feedback operators (or even zero) is desirable to optimize the expansion and simplification procedure of CPTs, which is of great interest in performing compatibility and verification checks.

Feedback-parallel translation. The *feedback-parallel translation* strategy (FP) first composes all components in parallel, and then connects outputs to inputs by applying feedback operations (with appropriate wiring where necessary). FP is illustrated in Fig.4a, for the DelaySum component of Fig.2b. Split models the junction after wire a .




Fig. 5: Diagram ConstDiv illustrating interesting behavior of incremental translation.

Applying FP on the DelaySum diagram yields the following CPT:

$$\begin{aligned} \text{DelaySum} = & \text{feedback}^3([f, c, a, e, s \rightsquigarrow f, e, c, s, a] \\ & \circ (\text{Add} \parallel \text{UnitDelay} \parallel \text{Split}) \circ [c, a, s', f, g \rightsquigarrow f, c, a, s', g]) \end{aligned}$$

where $\text{feedback}^3(\cdot) = \text{feedback}(\text{feedback}(\text{feedback}(\cdot)))$ denotes application of the feedback operator 3 times, on the variables f , c , and a , respectively (recall that feedback works only on one variable at a time, the first input and first output of the argument transformer). In order to apply feedback^3 to the parallel composition $\text{Add} \parallel \text{UnitDelay} \parallel \text{Split}$, we first have to reorder its inputs and outputs, such that the variables on which the feedbacks are applied come first in matching order. This is achieved by the *rerouting* transformers $[f, c, a, e, s \rightsquigarrow f, e, c, s, a]$ and $[c, a, s', f, g \rightsquigarrow f, c, a, s', g]$.

Incremental translation. The *incremental translation* strategy (IT) composes components one by one, after having ordered them in topological order according to the dependencies in the diagram.

The IT strategy is illustrated in Fig.4b. First, topological sorting yields the order $\text{Add}, \text{UnitDelay}, \text{Split}$. So IT first composes Add and UnitDelay . Since the two are connected with c , serial composition is applied, obtaining the CPT

$$\text{ICC1} = (\text{Add} \parallel \text{Id}) \circ \text{UnitDelay}$$

Next, IT composes ICC1 with Split . This requires both serial composition and feedback, and yields the final CPT:

$$\text{DelaySum} = \text{feedback}(\text{ICC1} \circ (\text{Split} \parallel \text{Id}))$$

It is worth noting that composing systems incrementally in this way might result in not the most natural compositions. For example, consider the diagram from Fig.5. The “natural” CPT for this diagram is probably: $(\text{Const}(1) \parallel \text{Const}(0)) \circ \text{Div} \circ \text{Split} \circ (\text{Scope} \parallel \text{Scope})$. Instead, IT generates the following CPT: $(\text{Const}(1) \parallel \text{Const}(0)) \circ \text{Div} \circ \text{Split} \circ (\text{Scope} \parallel \text{Id}) \circ (\text{Id} \parallel \text{Scope})$. More sophisticated methods may be developed to extract parallelism in the diagram and avoid redundant Id compositions like in the above CPT. This study is left for future work.

Feedbackless translation. As we shall see in §5, feedback is problematic for simplification. The reason is that general feedback does not generally preserve the functional predicate transformer property. It also introduces formulas with quantifiers, as explained in §3, and quantifier elimination is not always easy. We would like therefore to avoid the feedback operator in the generated CPTs as much as possible. The *feedbackless translation* strategy (NFBT) avoids feedback altogether. The key idea is that, since the diagram has no algebraic loops, we should be able to eliminate feedback and replace it with direct operations on current- and next-state variables, just like with basic blocks. In particular, we can *decompose* UnitDelay into two Id transformers, denoted Id_{ud1} and Id_{ud2} : Id_{ud1} computes the next state from the input, while Id_{ud2} computes the output from the current state.

Generally, we decompose all components having multiple outputs into several components having each a single output. For each new component we keep only the inputs they depend on, as shown in Fig.4c. Thus, the Split component from Fig.4b is also divided into two Id components, denoted $\text{Id}_{\text{split1}}$ and $\text{Id}_{\text{split2}}$.

Decomposing into components with single outputs allows to compute a separate CPT for each of the outputs. Then we take the parallel composition of these CPTs to form the CPT of the entire diagram.

On the example from Fig.4c, this translation proceeds as follows. It takes the component UD2 and its *targets* Split2 and Split1 , and replaces them by the new components $\text{UD2} \circ \text{Split2}$ and $\text{UD2} \circ \text{Split1}$. Next it takes the component $\text{UD2} \circ \text{Split1}$ and its target Add , and replaces them by $((\text{UD2} \circ \text{Split1}) \parallel \text{Id}) \circ \text{Add}$. In the final step it takes $((\text{UD2} \circ \text{Split1}) \parallel \text{Id}) \circ \text{Add}$ and its target UD1 and replaces them by $((\text{UD2} \circ$

$\text{Split1}) \parallel \text{Id}) \circ \text{Add} \circ \text{UD1}$. This procedure stops when all components have been connected. In general we obtain one component for every output of the systems, and for this example we obtain two components. In the end these components are composed in parallel to get the CPT. For this example we obtain:

$$\text{DelaySum} = [s, e \rightsquigarrow s, s, e] \circ \left((\text{Id}_{\text{ud2}} \circ \text{Id}_{\text{splt2}}) \parallel (((\text{Id}_{\text{ud2}} \circ \text{Id}_{\text{splt1}}) \parallel \text{Id}) \circ \text{Add} \circ \text{Id}_{\text{ud1}}) \right)$$

Because Id_{ud1} , Id_{ud2} , Id_{splt1} and Id_{splt2} are all Ids, and $\text{Id} \circ A = A \circ \text{Id} = A$ and $\text{Id} \parallel \text{Id} = \text{Id}$ (thanks to polymorphism), this CPT is reduced to $\text{DelaySum} = [s, e \rightsquigarrow s, s, e] \circ (\text{Id} \parallel \text{Add})$. The direct generation of this simpler CPT is possible, and part of future work.

5 Implementation and Evaluation

We implemented the translation algorithms presented in §4 into a prototype tool called the `SIMULINK2ISABELLE` compiler, downloadable from <http://users.ics.aalto.fi/iulia/sim2isa.shtml>. In this section we present the tool and report evaluation results on several case studies, including an industrial-grade benchmark by Toyota [11,12].

5.1 Toolset

The `SIMULINK2ISABELLE` compiler, written in Python, takes as input Simulink files in XML format and produces valid Isabelle theories that can be subjected to compatibility checking and verification. The compiler currently handles a sufficiently large subset of the Simulink libraries able to express industrial-grade models such as the Toyota benchmarks. Extending the compiler to handle even more basic blocks is ongoing work.

During the parsing and preprocessing phase of the input Simulink file, the tool performs a set of checks, including algebraic loop detection, unsupported blocks and/or block parameters, malformed blocks (e.g., a function block referring to a nonexistent input), etc., and issues possible warnings/errors.

The tool implements all three translation strategies, and takes two additional options: `flat` (flatten diagram) and `io` (intermediate outputs). Option `flat` flattens the hierarchy of the HBD and produces a single diagram consisting only of basic blocks (no subsystems), on which the translation is then applied. Note that FP and IT are by design compositional, i.e., they preserve the hierarchical structure of the model. However, using the `flat` option discards this feature. Option `io` generates and names all intermediate CPTs produced during the translation process. These names are then used in the CPT for the top-level system, to make it shorter and more readable. In addition, the intermediate CPTs can be expanded and simplified incrementally by Isabelle, and used in their simplified form when computing the CPT for the next level up. This generally results in more efficient simplification. Another benefit of producing intermediate CPTs is the detection of incompatibilities early during the simplification phase. Moreover, this indicates the group of components at fault and helps localize the error.

Isabelle is a proof assistant, but not a fully-automatic tool. We have therefore implemented in Isabelle a set of functions (keyword `simulink`) which result in the generated CPTs being expanded and simplified automatically. By *expansion* we mean replacing the serial, parallel, and feedback composition operators based on the Lemmas 1 and 2. By *simplification* we mean simplifying the resulting formulas, e.g., by eliminating quantifiers and internal variables. The goal is to obtain for the top-level system a single atomic predicate transformer which only refers to the external input, output, and state/next state variables of the system (and not to the internal wires of the diagram).

For instance, when executed on our running example (Fig.2) with the IT option, the tool produces the Isabelle code:

```
simulink DelaySum = feedback((Add || Id) o UnitDelay o (Split || Id))
simulink Counter = (Const(1) || Id) o DelaySum o (Scope || Id)
```

When executed in Isabelle, this code automatically generates the definitions (2) as well as the simplification theorems (3), and automatically proves these theorems.

As another example, when we run the tool on the example of Fig.5, we obtain the theorem `ConstDiv = {x : false}`, which states that the system has no legal inputs. This indicates incompatibility, due to performing a division by zero.

The generated `simulink` declarations preserve the name of the block instance they map, therefore enforcing the traceability of the translation.

	FP		IT				NFBT
	PHBD	FHBD	PHBD	FHBD	IO-PHBD	IO-FHBD	
\mathcal{L}_{cpt}	731	598	1317	1546	1325	1303	2455
\mathcal{N}_{cpt}	2	1	2	1	6	6	4
Readability	medium	medium	medium	medium	high	high	low
\mathcal{T}_{exp}	0.095	0.114	0.120	0.154	0.193	0.199	0.102
\mathcal{P}_{exp}	0.013	0.016	0.024	0.017	0.018	0.017	0.107
$\mathcal{N}_{\text{quant}}$	11	9	19	21	20	20	0
\mathcal{T}_{tot}	0.247	0.312	0.220	0.338	0.263	0.274	0.072
$\mathcal{P}_{\text{simp}}$	0.01	0.0	0.0	0.0	0.0	0.0	0

Table 1: Evaluation results for the running example (Fig.2).

	FP		IT				NFBT
	PHBD	FHBD	PHBD	FHBD	IO-PHBD	IO-FHBD	
\mathcal{L}_{cpt}	17191	15955	141989	167291	155871	189581	33372
\mathcal{N}_{cpt}	6	1	6	1	84	84	101
\mathcal{T}_{exp}	∞	∞	∞	∞	1809.184	∞	5.564
\mathcal{P}_{exp}	-	-	-	-	∞	-	27.017
$\mathcal{N}_{\text{quant}}$	-	-	-	-	-	-	0
\mathcal{T}_{tot}	-	-	-	-	3393.787	3892.872	14.089
$\mathcal{P}_{\text{simp}}$	-	-	-	-	1.756	2.795	1.140
$\mathcal{L}_{\text{simp}}$	-	-	-	-	41047	45761	21243

Table 2: Evaluation results for the FCS Simulink model [11,12].

5.2 Evaluation

We present evaluation results from two case studies: the running example (Fig.2) and the Fuel Control System (FCS) model described in [11,12]. FCS solves the problem of maintaining the ratio of air mass and injected fuel at the stoichiometric value [5], i.e., enough air is provided to completely burn the fuel in a car’s engine. This control problem has important implications on lowering pollution and improving engine performance. Three designs are presented in [11,12], modeled in Simulink, Hybrid I/O Automata [16] and Polynomial Hybrid I/O Automata [9]. We evaluate our approach on the Simulink model, available from <http://cps-vo.org/group/ARCH/benchmarks>. The model has a 3-level hierarchy with a total of 70 block instances (65 basic blocks and 5 subsystems), connected by 82 wires, of which 8 feedbacks.

We evaluate the three translation strategies with the options: FP without/with flattening (PHBD/FHBD), IT without/with flattening and without/with io option (IO), and NFBT. NFBT by construction generates intermediate outputs and does not preserve the structure of the hierarchy in the result, thus, its result is identical with/without the options. The results from the running example are shown in Table 1 and from the FCS model in Table 2. Table 3 shows results from the largest flat subsystem of the FCS model. This subsystem has 35 basic blocks and 47 wires.

The evaluation criteria are: (1) \mathcal{L}_{cpt} : length of the produced CPTs (# characters), (2) \mathcal{N}_{cpt} : number of generated CPTs, (3) readability, (4) \mathcal{T}_{exp} and \mathcal{P}_{exp} : times required for the expansion of CPTs and for the printing of the top-level formula, (5) $\mathcal{N}_{\text{quant}}$: number of quantifiers in the top-level formula, (6) \mathcal{T}_{tot} : total time needed for expansion and simplification, (7) $\mathcal{P}_{\text{simp}}$: time to print the simplified formula, (8) $\mathcal{L}_{\text{simp}}$: length of the simplified formula (# chars). All times are in seconds. The translation time itself (i.e., the time to process the Simulink file and generate the Isabelle file) is negligible for all three strategies (at most one fifth of a second) and therefore not reported. We present separately times to expand/simplify and times to print the formulas, since printing takes significant time in the Isabelle/ML framework.

Readability is subjective, of course, and our measures are relative. Readability is higher in the incremental strategy with io option, since the intermediate outputs allows to parse the result step by step. Readability is low in NFBT because this method decomposes blocks and does not preserve the hierarchy

	FP	IT		NFBT
		PHBD	IO-PHBD	
\mathcal{L}_{cpt}	13042	88059	85116	19258
\mathcal{N}_{cpt}	1	1	47	56
\mathcal{T}_{exp}	32.197	287.874	243.86	3.029
\mathcal{P}_{exp}	456.758	371.601	276.799	4.527
\mathcal{L}_{exp}	176229	176229	216674	209125
$\mathcal{N}_{\text{quant}}$	176	176	193	0
\mathcal{T}_{tot}	1689.427	1401.053	969.684	5.524
$\mathcal{P}_{\text{simp}}$	0.603	0.571	0.525	0.327
$\mathcal{L}_{\text{simp}}$	35115	17577	17577	17303

Table 3: Evaluation results for the largest subsystem of the FCS model.

of the original model. For this reason, and even though NFBT is superior in terms of efficiency, the other methods are still interesting to pursue and improve as part of future work.

The ∞ symbol denotes timeout after two hours of computation. The reason Isabelle fails to expand or print some formulas is that they become too large. Also, expansion involves more than just inlining, e.g., renaming variables appropriately due to quantifiers. We note that $\text{feedback}(\{p\} \circ [r])$ expands to a formula with length roughly equal to $(3 \cdot |p| + |r|) + (|p| + 2 \cdot |r|)$. Similarly, feedback introduces $3 + 3 \cdot \mathcal{N}_{\text{quant}}(p) + \mathcal{N}_{\text{quant}}(r)$ quantifiers in the precondition and $3 + \mathcal{N}_{\text{quant}}(p) + 2 \cdot \mathcal{N}_{\text{quant}}(r)$ in the relation. Note also that a feedback wire can transfer an array of n values, which is translated by our tool as n successive applications of feedback .

Observe (Table 2) that with IO-PHBD Isabelle manages to expand the CPTs (into an internal data structure) after about 30 mins, but not to print the expanded formula. However, the expanded formula can be simplified in less than 30 mins and the simplified formula can be printed in less than 2 secs. In the case of NFBT, everything (expansion+simplification+printing) takes less than a sec. The final simplified formulas are in all cases quantifier-free.

6 Related Work

A plethora of work exists on translating Simulink models to various target languages, for verification purposes or for code-generation purposes. Although some of these works have ultimate goals similar to ours (e.g., verification), to the best of our knowledge, none of them studies the same problem that we study in this paper: the different ways to translate HBDs into a compositional algebra with serial, parallel, and feedback composition operators.

Primarily focusing on verification and targeting discrete-time fragments of Simulink, existing works describe translations to BIP [21], NuSMV [18], or Lustre [23]. Other works study transformation of continuous-time Simulink to Timed Interval Calculus [4], Function Blocks [10], I/O Extended Finite Automata [24], or Hybrid CSP [25]. The Stateflow module of Simulink, which allows to model hierarchical state machines, has been the subject of translation to hybrid automata [1,17].

Contract-based frameworks for Simulink are described in [3,20]. Although these works are related to ours, they do not solve the same problem as explained above. [3] uses pre/post-conditions as contracts for discrete-time Simulink blocks and SDF graphs [13] to represent Simulink diagrams. Then sequential code is generated from the SDF graph, and the code is verified using traditional refinement-based techniques [2] and the Z3 SMT solver [6]. In [20] Simulink blocks are annotated with rich types (separate constraints on inputs and outputs, but no relations between inputs and outputs which is possible in our framework). Then the SimCheck tool extracts verification conditions from the Simulink model and the annotations, and submits them to the Yices SMT solver [8] for verification.

Modular code generation methods for Simulink models are described in [15,14]. The main technical problem solved there is how to cluster subsystems in as few clusters as possible without introducing false input-output dependencies.

7 Conclusion

In this paper we tackle the problem of translating Simulink diagrams (and more generally, hierarchical block diagrams) into a compositional framework: an algebra of atomic systems composed in series, parallel, or feedback. The challenge comes from the fact that there are many ways to translate a diagram into a term in this algebra, with different tradeoffs. Three translation strategies are proposed and implemented in a compiler that generates verifiable Isabelle code. These strategies are illustrated on a simple example and evaluated on a real-life Simulink model from Toyota.

Two intuitive translation strategies are first considered: the feedback-parallel translation and the incremental translation. These strategies have the benefit of producing relatively readable composite predicate transformers, which can be traced back to the structure of the original diagram. Unfortunately, expansion and simplification of the CPTs resulting from these translations do not scale well in Isabelle. The culprits are the length of the predicate transformers and the feedback composition operators that introduce quantifiers and are resource consuming. A third strategy – feedbackless – was designed to handle this scalability issue by avoiding feedback altogether. The measurements obtained show that the feedbackless method is by far the most efficient on all examples we tried. The drawback of the feedbackless method is that it decomposes the blocks of the original diagram, which renders this blocks difficult to identify in the produced CPTs. In consequence, the resulting expressions are harder to understand and diagnose in case of errors.

Future work directions include: (1) studying other translation strategies; (2) improving the automated simplification methods within Isabelle or other solvers; (3) extending the tool to handle a larger subset of Simulink; (4) extending the tool with automatic verification methods (model-checking requirements against the top-level CPT); and (5) extending the tool with fault localization methods whenever the compatibility or verification checks fail.

Another interesting question regards the semantical equivalence of the CPTs generated by the various strategies and options. We proved in Isabelle that the final simplified formulas are equivalent, for the running example as well as for several other examples and subsystems of the FCS model, including the largest subsystem reported in Table 3. It is future work to prove a meta-theorem that states that these equivalences hold by construction on any algebraic-loop-free example.

References

1. A. Agrawal, G. Simon, and G. Karsai. Semantic Translation of Simulink/Stateflow Models to Hybrid Automata Using Graph Transformations. *Electronic Notes in Theoretical Computer Science*, 109:43 – 56, 2004. Proceedings of the Workshop on Graph Transformation and Visual Modelling Techniques (GT-VMT 2004).
2. R.-J. Back and J. von Wright. *Refinement Calculus. A systematic Introduction*. Springer, 1998.
3. P. Boström. Contract-Based Verification of Simulink Models. In S. Qin and Z. Qiu, editors, *Formal Methods and Software Engineering*, volume 6991 of *Lecture Notes in Computer Science*, pages 291–306. Springer Berlin Heidelberg, 2011.
4. C. Chen, J. S. Dong, and J. Sun. A formal framework for modeling and validating Simulink diagrams. *Formal Aspects of Computing*, 21(5):451–483, 2009.
5. J. A. Cook, J. Sun, J. H. Buckland, I. V. Kolmanovsky, H. Peng, and J. W. Grizzle. Automotive Powertrain Control – A Survey. *Asian Journal of Control*, 8(3):237–260, 2006.
6. L. De Moura and N. Bjørner. Z3: An Efficient SMT Solver. In *Proceedings of the Theory and Practice of Software, 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, TACAS’08/ETAPS’08, pages 337–340, Berlin, Heidelberg, 2008. Springer-Verlag.
7. E. Dijkstra. Guarded commands, nondeterminacy and formal derivation of programs. *Comm. ACM*, 18(8):453–457, 1975.
8. B. Dutertre and L. de Moura. The Yices SMT solver. Technical report, SRI International, 2006.
9. G. Frehse, Z. Han, and B. Krogh. Assume-guarantee reasoning for hybrid I/O-automata by overapproximation of continuous interaction. In *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, volume 1, pages 479–484 Vol.1, Dec 2004.
10. C. han Yang and V. Vyatkin. Transformation of Simulink models to IEC 61499 Function Blocks for verification of distributed control systems . *Control Engineering Practice* , 20(12):1259 – 1269, 2012.
11. X. Jin, J. Deshmukh, J. Kapinski, K. Ueda, and K. Butts. Benchmarks for model transformations and conformance checking. In *1st Intl. Workshop on Applied Verification for Continuous and Hybrid Systems (ARCH)*, 2014.

12. X. Jin, J. V. Deshmukh, J. Kapinski, K. Ueda, and K. Butts. Powertrain Control Verification Benchmark. In *Proceedings of the 17th International Conference on Hybrid Systems: Computation and Control*, HSCC'14, pages 253–262. ACM, 2014.
13. E. Lee and D. Messerschmitt. Synchronous data flow. *Proceedings of the IEEE*, 75(9):1235–1245, 1987.
14. R. Lubliner, C. Szegedy, and S. Tripakis. Modular code generation from synchronous block diagrams – modularity vs. code size. In *36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'09)*, pages 78–89. ACM, Jan. 2009.
15. R. Lubliner and S. Tripakis. Modularity vs. reusability: Code generation from synchronous block diagrams. In *Design, Automation, and Test in Europe (DATE'08)*, pages 1504–1509. ACM, Mar. 2008.
16. N. Lynch, R. Segala, and F. Vaandrager. Hybrid I/O Automata. *Inf. Comput.*, 185(1):105–157, Aug. 2003.
17. K. Manamcheri, S. Mitra, S. Bak, and M. Caccamo. A Step Towards Verification and Synthesis from Simulink/Stateflow Models. In *Proceedings of the 14th International Conference on Hybrid Systems: Computation and Control*, HSCC '11, pages 317–318, New York, NY, USA, 2011. ACM.
18. B. Meenakshi, A. Bhatnagar, and S. Roy. Tool for Translating Simulink Models into Input Language of a Model Checker. In Z. Liu and J. He, editors, *Formal Methods and Software Engineering*, volume 4260 of *Lecture Notes in Computer Science*, pages 606–620. Springer Berlin Heidelberg, 2006.
19. V. Preoteasa and S. Tripakis. Refinement calculus of reactive systems. In *Embedded Software (EMSOFT), 2014 International Conference on*, pages 1–10, Oct 2014.
20. P. Roy and N. Shankar. SimCheck: a contract type system for Simulink. *Innovations in Systems and Software Engineering*, 7(2):73–83, 2011.
21. V. Sfyrla, G. Tsiligiannis, I. Safaka, M. Bozga, and J. Sifakis. Compositional translation of simulink models into synchronous BIP. In *Industrial Embedded Systems (SIES), 2010 International Symposium on*, pages 217–220, July 2010.
22. S. Tripakis, B. Lickly, T. A. Henzinger, and E. A. Lee. A Theory of Synchronous Relational Interfaces. *ACM Trans. Program. Lang. Syst.*, 33(4):14:1–14:41, July 2011.
23. S. Tripakis, C. Sofronis, P. Caspi, and A. Curic. Translating Discrete-time Simulink to Lustre. *ACM Trans. Embed. Comput. Syst.*, 4(4):779–818, Nov. 2005.
24. C. Zhou and R. Kumar. Semantic Translation of Simulink Diagrams to Input/Output Extended Finite Automata. *Discrete Event Dynamic Systems*, 22(2):223–247, 2012.
25. L. Zou, N. Zhany, S. Wang, M. Franzle, and S. Qin. Verifying Simulink diagrams via a Hybrid Hoare Logic Prover. In *Embedded Software (EMSOFT), 2013 Proceedings of the International Conference on*, pages 1–10, Sept 2013.